TABLE OF CONTENTS

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
BEFORE THE BOARD OF PATENT APPEALS AND INTERFERENCES

In re Application of	:	Customer Number: 46320
Jeffrey CHASE et al.	:	Confirmation Number: 2068
Application No.: 10/733,996	:	Group Art Unit: 2157
Filed: December 11, 2003	:	Examiner: B. Burgess

For: AUTONOMIC SELECTION OF A REQUEST ROUTING POLICY BASED UPON CACHE EFFECTIVENESS

APPEAL BRIEF

Mail Stop Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

Sir:

This Appeal Brief is submitted in support of the Notice of Appeal filed July 11, 2008, wherein Appellants appeal from the Examiner's rejection of claims 1-14.

I. REAL PARTY IN INTEREST

This application is assigned to IBM Corporation by assignment recorded on December 11, 2003, at Reel 014802, Frame 0372.

II. RELATED APPEALS AND INTERFERENCES

Appellants are unaware of any related appeals and interferences.

III. STATUS OF CLAIMS

Claims 1-14 are pending and two-times rejected in this Application. It is from the multiple rejections of claims 1-14 that this Appeal is taken.

IV. STATUS OF AMENDMENTS

The claims have not been amended subsequent to the imposition of the Second and Final Office Action dated May 29, 2008 (hereinafter the Second Office Action).

V. SUMMARY OF CLAIMED SUBJECT MATTER

Referring to Figure 1 and also to independent claim 1, an autonomic request routing policy selection system is disclosed. The request routing policy selection system includes a plurality of pre-configured request routing policies 120, a data store 140, and a routing policy selector 200. The data store 140 is of cache metrics for the pre-configured request routing policies 120 (lines 1-6 of paragraph [0022]). The routing policy selector 200 is configured for communicative linkage to a server cluster 110 comprising a plurality of servers and programmed to select a particular one of the request routing policies 120 for use in routing content requests 130 in the server cluster 110 based upon the cache metrics 140 (lines 1-9 of paragraph [0021]). The routing policy selector 200 further comprises a coupling to the routing policies 120 and the data store of cache metrics 140 (lines 7-9 of paragraph [0021]).

Referring to Figure 2 and also to independent claim 7, an autonomic request routing policy selection method is disclosed. In block 210, a contemporary trace footprint experienced by a coupled server cluster is identified (lines 3-4 of paragraph [0023]). In block 220, a cache allocation for the coupled server cluster is identified (lines 4-5 of paragraph [0024]). In block

230, at least two sets of hit rate metrics are retrieved, and each set of metrics correspond to a particular routing policy (lines 5-9 of paragraph [0023]). In block 240, the hit rate metrics are compared based upon the identified trace footprint and the identified cache allocation to determine a preferred routing policy (lines 1-8 of paragraph [0024]). In block 250, the preferred routing policy is selected for use in routing content requests to the server cluster (lines 8-9 of paragraph [0024]).

Referring to Figure 2 and also to independent claim 11, a machine readable storage having stored thereon a computer program for autonomic request routing policy selection is disclosed. The computer program comprises a routine set of instructions which when executed by the machine cause the machine to perform the following steps. In block 210, a contemporary trace footprint experienced by a coupled server cluster is identified (lines 3-4 of paragraph [0023]). In block 220, a cache allocation for the coupled server cluster is identified (lines 4-5 of paragraph [0024]). In block 230, at least two sets of hit rate metrics are retrieved, and each set of metrics correspond to a particular routing policy (lines 5-9 of paragraph [0023]). In block 240, the hit rate metrics are compared based upon the identified trace footprint and the identified cache allocation to determine a preferred routing policy (lines 1-8 of paragraph [0024]). In block 250, the preferred routing policy is selected for use in routing content requests to the server cluster (lines 8-9 of paragraph [0024]).

VI. GROUNDS OF REJECTION TO BE REVIEWED ON APPEAL

1. Claim 1 was rejected under 35 U.S.C. § 102 for anticipation based upon Mangipudi et al., U.S. Patent Publication No. 2004/0162901 (hereinafter Mangipudi);

2. Claims 2-6 were rejected under 35 U.S.C. § 103 for obviousness based upon Mangipudi in view of Appellants' Admitted Prior Art (hereinafter the Admitted Prior Art); and

3. Claims 7-14 were rejected under 35 U.S.C. § 103 for obviousness based upon Mangipudi in view of Yang, U.S. Patent Publication No. 2002/0199014 (hereinafter Yang).

VII. ARGUMENT

As is evident from Appellants' previously-presented comments during prosecution of the present Application and from Appellants' comments below, there are questions as to how the limitations in the claims correspond to features in the applied prior art. In this regard, reference is made to M.P.E.P. § 1207.02, entitled "Contents of Examiner's Answer." Specifically, the following is stated:

> (A) CONTENT REQUIREMENTS FOR EXAMINER'S ANSWER. The examiner's answer is required to include, under appropriate headings, in the order indicated, the following items:
>
> ...
>
> (9)(c) For each rejection under 35 U.S.C. 102 or 103 where there are questions as to how limitations in the claims correspond to features in the prior art even after the examiner complies with the requirements of paragraphs (c) and (d) of this section, the examiner must compare at least one of the rejected claims feature by feature with the prior art relied on in the rejection. The comparison must align the language of the claim side-by-side with a reference to the specific page, line number, drawing reference number, and quotation from the prior art, as appropriate. (emphasis added)

Therefore, if the Examiner is to maintain the present rejections and intends to file an Examiner's Answer, the Examiner is required to include the aforementioned section in the Examiner's Answer.

Appellants have compared the statement of the rejection found on pages 2-14 of the Second Office Action with the statement of the rejection found on pages 2-15 of the First Office Action. Upon making this comparison, Appellants have been unable to discover any substantial differences between the respective statements of the rejection. As such, Appellants proceed on the basis that the Examiner's sole response to Appellants' First Response dated February 13, 2008 (hereinafter the First Response) is found on pages 14 and 15 of the Second Office Action in the section entitled "Response to Arguments."

THE REJECTION OF CLAIM 1 UNDER 35 U.S.C. § 102 FOR ANTICIPATION BASED UPON MANGIPUDI

For convenience of the Honorable Board in addressing the rejections, independent claim 1 stands or falls alone.

The factual determination of anticipation under 35 U.S.C. § 102 requires the identical disclosure, either explicitly or inherently, of each element of a claimed invention in a single reference.[1] Moreover, the anticipating prior art reference must describe the recited invention with sufficient clarity and detail to establish that the claimed limitations existed in the prior art and that such existence would be recognized by one having ordinary skill in the art.[2] As part of this analysis, the Examiner must (a) identify the elements of the claims, (b) determine the meaning of the elements in light of the specification and prosecution history, and (c) identify corresponding elements disclosed in the allegedly anticipating reference.[3] This burden has not been met.

Claim 1, in part, recites the following limitation:

> a routing policy selector configured for communicative linkage to a server cluster comprising a plurality of servers and programmed to select a particular one of said request routing policies for use in routing content requests in said server cluster based upon said cache metrics, said routing policy selector further

[1] In re Rijckaert, 9 F.3d 1531, 28 USPQ2d 1955 (Fed. Cir. 1993); Richardson v. Suzuki Motor Co., 868 F.2d 1226, 1236, 9 USPQ2d 1913, 1920 (Fed. Cir. 1989); Perkin-Elmer Corp. v. Computervision Corp., 732 F.2d 888, 894, 221 USPQ 669, 673 (Fed. Cir. 1984).
[2] See In re Spada, 911 F.2d 705, 708, 15 USPQ 1655, 1657 (Fed. Cir. 1990); Diversitech Corp. v. Century Steps, Inc., 850 F.2d 675, 678, 7 USPQ2d 1315, 1317 (Fed. Cir. 1988).
[3] Lindermann Maschinenfabrik GMBH v. American Hoist & Derrick Co., 730 F.2d 1452, 221 USPQ 481 (Fed. Cir. 1984).

comprising a coupling to said routing policies and said data store of cache metrics.

In the paragraph spanning pages 3 and 4 of the First Office Action, the Examiner cited paragraphs [0018], [0020], [0037], and [0039] of Mangipudi and asserted the following:

> Man teaches a routing host (router) having a policy engine used to make routing decisions using parameters and defined policies. The router is connected to server host referred to as cluster. It also keeps a database with policies and parameter used to make routing decisions.

Absent from both the Examiner's analysis and the cited passages within Mangipudi, however, is a discussion of selecting a request routing policies based upon cache metrics. Paragraph [0018] of Mangipudi describes routing and load balancing, but this passage discusses neither selecting a request routing policy nor cache metrics. Paragraph [0020] describes that routing decisions are based upon information in a Management Information Base (MIB); however, this passage is silent as to selecting a particular request routing policy based upon cache metrics. Paragraph [0037] describes that requests are routed based on policies, but this passage does not describe how the policies are selected or contain any mention of cache metrics. Paragraph [0039] is a discussion of the general mechanics involved in routing requests but is silent as to selecting a request routing policy or cache metrics. Therefore, Appellants submit that the Examiner has failed to establish that Mangipudi identically discloses the claimed routing policy selector.

For above-described reasons, the Examiner has failed to establish that Mangipudi identically discloses the claimed invention, as recited in claim 1, within the meaning of 35 U.S.C. § 102. The aforementioned arguments were previously presented on pages 2 and 3 of the First Response, and the Examiner responded to these arguments on page 15 of the Second Office Action

as follows:

> Mangipudi teaches a policy engine used to select load distribution algorithms for finding the best route for traffic to backend servers (paragraphs [0018-0020]). Attributes or metrics used in the selecting process are response time, request, transaction type, protocol type, file size, online/offline status, CPU utilization, etc. (paragraph [0019]).

To be clear, as claimed, the selection of a particular one of the request routing policies is based upon cache metrics.

Regarding the "load distribution algorithms" referred to by the Examiner, paragraph [0018] of Mangipudi states:

> The APE or policy engine in conjunction with the router then intelligently distributes incoming traffic to the most available and/or efficient server within each class or "cluster," by using one or more of a plurality of selectable load distribution algorithms for the class/cluster, including: weighted percentage load balancing; round robin load balancing; CPU availability load balancing; probabilistic load balancing and least connections load balancing. Thus each back-end server in communication with the router is subject to a selectable one of a plurality of load balancing algorithms so that traffic is routed to the plurality of back-end servers, as a function of class, in a manner that maintains consistent response times and service level commitments even with increases in traffic and processing loads.

Notably, Mangipudi does not teach the selection of one of the load distribution algorithms (i.e., allegedly corresponding to the claimed request routing policies) is based upon cache metrics. Instead, referring to paragraph [0020], Mangipudi teaches:

> These parameters reported to the router are available via a Management Information Base (MIB) kept by the policy engine on the router. The information is also made available by the policy engine to application layer programs (i.e. on an NT platform via performance monitoring registers, "perfmon registers"). The policy engine uses this information, in conjunction with the router, in making load distribution decisions. The policy engine in conjunction with the router uses the information/parameters to determine the configuration of the class and cluster(s), as well as in making load distribution decisions.

As discussed therein, Mangipudi teaches make routing decisions based upon information reported to the router and stored in a MIB (see paragraph [0019] for a discussion of the information reported to the router). Making routing decisions based upon certain information is different than making a selection decision as to a particular request routing policy to employ.

Moreover, the Examiner has still failed to establish that Mangipudi teaches that the information used to make the decision is based upon cache metrics.

THE REJECTION OF CLAIMS 2-6 UNDER 35 U.S.C. § 103 FOR OBVIOUSNESS BASED UPON MANGIPUDI IN VIEW OF THE ADMITTED PRIOR ART

For convenience of the Honorable Board in addressing the rejections, claims 2-6 stand or fall together with independent claim 1.

Claims 2-6 depend from independent claim 1, and Appellants incorporate herein the arguments previously advanced in traversing the imposed rejection of claim 1 under 35 U.S.C. § 102 for anticipation based upon Mangipudi. The secondary reference to the Admitted Prior Art does not cure the argued deficiencies of the prior rejection. Accordingly, even if one having ordinary skill in the art were impelled to combine the applied prior art, the claimed invention would not result. Appellants, therefore, respectfully submit that the imposed rejection of claim 1 under 35 U.S.C. § 103 for obviousness based upon Mangipudi in view of the Admitted Prior Art is not viable.

THE REJECTION OF CLAIMS 7-14 UNDER 35 U.S.C. § 103 FOR OBVIOUSNESS BASED UPON MANGIPUDI IN VIEW OF YANG

For convenience of the Honorable Board in addressing the rejections, claims 8-14 stand or fall together with independent claim 7.

Claim 7, in part, recites the following limitation:

> retrieving at least two sets of hit rate metrics, each set of metrics corresponding to a particular routing policy.

In the paragraph spanning pages 7 and 8 of the First Office Action, the Examiner cited paragraph [0019], [0020], [0026], and [0038] of Mangipudi and asserted the following:

> Man teaches reporting to the policy engine total hits second, CPU utilization, response times of servers, URL/content availability. These are all sets of hit rate metrics. The policy engine uses these parameters in policies to route requests. Policies are class of service (CoS) and service level agreement (SLA).

At the outset, Appellants disagree with the Examiner's characterization of the claimed "sets of hit rate metrics." CPU utilization, response times of servers, URL/content availability are not sets of hit rate metrics, as claimed. Moreover, the Examiner has failed to provide any factual support to (or analysis for) the notion that one having ordinary skill in the art, given the broadest, reasonable interpretation of the phrase "sets of hit rate metrics" consistent with Appellants' specification, would consider CPU utilization, response times of servers, URL/content availability as being comparable to the claimed sets of hit rate metrics.

Appellants also note that the Examiner has failed to establish that Mangipudi disclose the claimed "each set of metrics corresponding to a particular routing policy." Of the Examiner's cited passages, the only discussion of hit rates is found in paragraph [0019]. However, this paragraph merely describes the information that is being collected "for each virtual site and for each web farm" and does not associate a set of metrics with a particular routing policy. Accordingly, even if one having ordinary skill in the art would have been impelled to combine the applied prior art, the claimed invention would not result from the Examiner's combination.

In response to the above-reproduced arguments, the Examiner asserted the following on page 15 of the Second Office Action:

> Applicant's specification describes hit rate metrics as specific workloads, particular cache size, available cache memory in the server cluster (Specification page 10, paragraph [0018], page 11, paragraph [0022]).
> Mangipudi teaches these metrics in paragraphs [0019, 0055].

At the outset, Appellants note that the Examiner has mischaracterized Appellants' specification. Although cache size, specific workloads, and available cache memory may affect a hit rate metric, these characteristics, by themselves, are not hit rate metrics.

Appellants also note that the Examiner's analysis fails to consider the actual language of the claims. The Examiner asserted that Mangipudi teaches "retrieving at least two sets of hit rate metrics, each set of metrics corresponding to a particular routing policy." Thus, the Examiner must establish that Mangipudi teaches a first set of hit rate metrics (i.e., a first set that includes a plurality of metrics) and a second set of hit rate metrics (i.e., a second set that includes a plurality of metrics). Moreover, the Examiner must establish that Mangipudi teaches that each of the two sets correspond to a particular routing policy. The Examiner's analysis however, has failed to identify either of the first or second sets and what metrics constitute the first and second sets. Moreover, the Examiner's analysis has failed to identify the particular routing policies with which the first and second set respectively correspond. Thus, the Examiner has still failed to properly identify the scope and content of Mangipudi.

Conclusion

Based upon the foregoing, Appellants respectfully submit that the Examiner's rejections under 35 U.S.C. §§ 102, 103 based upon the applied prior art is not viable. Appellants, therefore,

respectfully solicit the Honorable Board to reverse the Examiner's rejections under 35 U.S.C. §§ 102, 103.

To the extent necessary, a petition for an extension of time under 37 C.F.R. § 1.136 is hereby made. Please charge any shortage in fees due under 37 C.F.R. §§ 1.17, 41.20, and in connection with the filing of this paper, including extension of time fees, to Deposit Account 09-0461, and please credit any excess fees to such deposit account.

Date: July 11, 2008

Respectfully submitted,

/Scott D. Paul/
Scott D. Paul
Registration No. 42,984
Steven M. Greenberg
Registration No. 44,725
Phone: (561) 922-3845
CUSTOMER NUMBER 46320

VIII. CLAIMS APPENDIX

1. An autonomic request routing policy selection system comprising:

a plurality of pre-configured request routing policies;

a data store of cache metrics for said pre-configured request routing policies; and,

a routing policy selector configured for communicative linkage to a server cluster comprising a plurality of servers and programmed to select a particular one of said request routing policies for use in routing content requests in said server cluster based upon said cache metrics, said routing policy selector further comprising a coupling to said routing policies and said data store of cache metrics.

2. The system of claim 1, wherein said pre-configured request routing policies comprise a Layer 4 request routing policy and a Layer 7 request routing policy.

3. The system of claim 2, wherein said Layer 4 request routing policy comprises a server load balancing type policy.

4. The system of claim 2, wherein said Layer 7 request routing policy comprises a content localizing type policy.

5. The system of claim 4, wherein said content localizing type policy comprises a uniform resource locator (URL) hashing policy.

6. The system of claim 1, wherein said cache metrics comprises a plurality of Zipf-like analyses based upon different selected alpha values for different workloads imposed upon said server cluster according to different ones of said request routing policies.

7. An autonomic request routing policy selection method comprising the steps of:

identifying a contemporary trace footprint experienced by a coupled server cluster;

identifying a cache allocation for said coupled server cluster;

retrieving at least two sets of hit rate metrics, each set of metrics corresponding to a particular routing policy;

comparing said hit rate metrics based upon said identified trace footprint and said identified cache allocation to determine a preferred routing policy; and,

selecting said preferred routing policy for use in routing content requests to said server cluster.

8. The method of claim 7, further comprising the steps of:

computing with said hit rate metrics, an optimal server cluster configuration for said preferred routing policy; and,

provisioning an optimal number of servers in said server cluster based upon said computed optimal server cluster configuration.

9. The method of claim 7, wherein said selecting step comprises the step of selecting a server load balancing type routing policy when said identified cache allocation approaches in value said identified trace footprint.

10. The method of claim 7, wherein said selecting step comprises the step of selecting a content localizing type routing policy when either said identified cache allocation is small, or when said trace footprint is large.

11. A machine readable storage having stored thereon a computer program for autonomic request routing policy selection, the computer program comprising a routine set of instructions which when executed by the machine cause the machine to perform the steps of:

identifying a contemporary trace footprint experienced by a coupled server cluster;

identifying a cache allocation for said coupled server cluster;

retrieving at least two sets of hit rate metrics, each set of metrics corresponding to a particular routing policy;

comparing said hit rate metrics based upon said identified trace footprint and said identified cache allocation to determine a preferred routing policy; and,

selecting said preferred routing policy for use in routing content requests to said server cluster.

12. The machine readable storage of claim 11, further comprising the steps of:

computing with said hit rate metrics, an optimal server cluster configuration for said preferred routing policy; and,

provisioning an optimal number of servers in said server cluster based upon said computed optimal server cluster configuration.

13. The machine readable storage of claim 11, wherein said selecting step comprises the step of selecting a server load balancing type routing policy when said identified cache allocation approaches in value said identified trace footprint.

14. The machine readable storage of claim 11, wherein said selecting step comprises the step of selecting a content localizing type routing policy when either said identified cache allocation is small, or when said trace footprint is large.

IX. EVIDENCE APPENDIX

No evidence submitted pursuant to 37 C.F.R. §§ 1.130, 1.131, or 1.132 of this title or of any other evidence entered by the Examiner has been relied upon by Appellants in this Appeal, and thus no evidence is attached hereto.

X. RELATED PROCEEDINGS APPENDIX

Since Appellants are unaware of any related appeals and interferences, no decision rendered by a court or the Board is attached hereto.